

1. In section 1.1.1.4 the requirement states, "The systems shall be able to process metadata for a minimum of 20,000 events per second per instance"

However, in section 1.1.1.8, it also states, "The systems shall be able to immediately handle metadata for 60,000 events per second at each location. Proposed architecture without major modification shall be scalable to 120,000 events per second."

How is the VA using the term instance here versus location?

Answer: "Instance" would be a single device, so you may have 3 instances in a single "location" that would be capable of processing 60,000 events per second.

2. We have, we believe, arrived at a one rack solution. However, we do not think it will be technically feasible to remain under 15,000 BTU per hour as required in section 1.1.1.23. How wooden is this restriction? We realize the VA has stated that they only have one rack at each TIC available. However, could the VA or the contractor pay for the additional BTU per hour in the datacenter? Would this be an acceptable solution?

Answer: No. 15,000 BTU is half of what we're currently allowed to use.

3. While we realize that the VA has a limitation of one rack in its TICs, could the vendor purchase additional racks of space as part of the contract price if they felt a better technical solution could be arrived at with more than one rack?

Answer: No, we're looking for the best technical solution that fits within our required constraints.

4. Is it the Government's intent to provide a sole source justification for IBM product on this request? The specifications are written directly to IBM product with little to no opportunity for open market responses.

Answer: No, the Government intends to compete this requirement. All potential offerors are encouraged to provide a description of their technical solution. Additionally, all potential offerors are free to provide feedback on the requirement as described in the draft Performance Work Statement and System Technical Requirements document. VA will duly consider your feedback.

5. Is there a technical requirement for the SIEM to pull logs from SPLUNK which is a significant investment VA has already made and had verbally shared interest in keeping?

Answer: Only metadata will go from Splunk to the SIEM. Splunk will remain as the log manager. Solution should be dual homed to allow for raw logs to be sent to SPLUNK and metadata to be sent to SIEM by collectors/agents.

6. In attachment 001 Section 5.2 General SIEM Requirements, it states "The Contractor shall provide a Test lab system that will reside in Martinsburg, WV and be configured to replicate the centralized SIEM system in the TIC Gateway production environment (the Lab will replicate only one TIC gateway). This shall include the same device types, model(s), and mirror the HA configuration implemented in the TIC Gateway production environment. The test lab is isolated from, and has no connectivity to, the production environment. The Test lab system shall have the ability to process 30,000 events per second" but in Attachment 003 1.1.1 System Technical Requirements [Requirement] 8, "The systems shall be able to immediately handle metadata for 60,000 events per second at each location. Proposed architecture without major modification shall be scalable to 120,000 events per second."

In order to provide an accurate response to PWS requirements, we need the actual EPS for the Test and Production environments. Please clarify, is the test environment is only 50% of production and how that is considered a mirror of the anticipated production environment?

Answer: The VA has four TICs. Each shall have high availability (HA). The design of the HA will vary depending on the solution described in each potential offeror's RFI response. The test lab shall be a replica of one of the TICs, to include full functionality. In other words, if there are five elements with HA in each TIC and a console in Martinsburg with HA to another console in Hines, the lab shall have those five elements as well as consoles to replicate all SIEM capabilities including HA design. If, however, the five elements are designed to meet the prescribed load/ capacity by leveraging 10 devices, the lab would then only have 5 devices. This is a cost savings plan. VA cannot replicate in the lab the volumes of traffic that would be experienced in production. Consequently there is no need to have full EPS-capable test environment. However, the test lab implementation shall be a technical representation of the complexity (not an exact mirror) of the production implementation, particularly if multiple devices are needed to meet EPS requirements. VA requires a high confidence in testing new rules and other configurations in the lab before implementation into operational systems. The product shall be configured to meet the maximum EPS stated in the PWS and System Technical Requirement document.

7. Will the VA accept the use of QRadar's native HA capability (which the VA uses on its existing QRadar systems) as meeting the definition of Active-Active?"

Answer: Yes, as long as all requirements are met